working paper
department
of economics

AN EXAMPLE OF NON-MARTINGALE LEARNING

Lones Smith
Peter Sorensen

94-25   July 1994

massachusetts
institute of
technology

50 memorial drive
cambridge, mass. 02139

AN EXAMPLE OF NON-MARTINGALE LEARNING

Lones Smith
Peter Sorensen

94-25   **July 1994**

An Example of Non-Martingale Learning.[*][†]

Lones Smith[‡]

and

Peter Sørensen[§]

Department of Economics
Massachusetts Institute of Technology
Cambridge, MA 02139

July 22, 1994

Abstract

This paper abandons the universal assumption in models of observational learning that individuals can perfectly observe the order of past decisions. Only a worked example so far, we show that complete learning obtains under the same conditions as before. We try to underscore that the resulting analysis in no way exploits any martingales — which have been the lifeblood of optimal learning theory until now.

1. INTRODUCTION

Consider the following canonical model: A sequence of individuals each must make a once-in-a-lifetime binary decision, with common payoff either 1 (in the high state) or -1 (in the low state) with equal chance. Everyone must decide on the basis of his own private signal and the knowledge of what all predecessors have done. Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) (henceforth, BHW) independently noticed that 'herding' may arise in the above observational learning model: There is a positive chance that everyone eventually settles upon the less profitable decision. Smith and Sorensen (1994) (hereafter SS) have recently shown, among other things, that the 'herding' results of Banerjee (1992) and BHW obtain precisely when the signals possessed by the individuals have bounded informativeness. Otherwise, barring other innovations in the model, complete learning obtains.

A crucial and perhaps untenable assumption in the whole 'herding' literature is the *perfect* observability of history, and in particular, knowledge of the exact order of all decisions: A modelling assumption to be sure, but how important is it? Smith (1991) thus posited that entrants can only observe the number of individuals choosing each option, and nothing else.[1] In this discrete player model,[2] we conjecture (and have almost proved) that complete learning arises under the same assumption of 'unbounded signals' as in SS. Knowing the exact order of all prior decisions is thus not asymptotically critical.

But this variation on the standard observational learning model in itself does not justify a heavy investment of effort, and we have no intention of making a curious commentary on the herding literature. Rather, we are hoping to push the understanding of how *social learning* occurs in whole new theoretical direction. For unlike all other 'herding' papers, the entire single-person learning literature (cf. Easley and Kiefer (1988)), and the rational expectations literature (cf. Bray and Kreps (1987)), what really drives convergence here is no longer the Martingale Convergence Theorem. For while individual 2 will always be able to infer the action of his predecessor, if he then takes the opposite action of individual 1, then individual 3 will not be able to infer who did what: In this sense, information is both gained and lost, and thus not refined, over time. Thus, in this worked example, where the analytics are especially clean, we try to show that there is no (non-trivial) conditional or unconditional martingale to found. While we wish to solve the model, rather than explain how *not* solve it, we do feel defensive on this point. We do discover and exploit a supermartingale, but this only turns out to be the more modest ingredient in a very complex picture. We then prove that full learning obtains in the limit. We have some simulations of the dynamics of the above learning model that attest to the nontrivial nature of what is going on here.

We eventually wish to use this result to develop a more encompassing economic model of learning, in which information is both gained and lost.

---

[1]Smith (1991) embellished the analysis by assuming individuals played 'two-armed bandits' following their observation of history. But the resulting 'error persistence' (incomplete learning) can just as easily be proxied by the possibility that individuals sometimes err in their one-shot decisions.

[2]The more tractable continuum player analysis in Smith (1991) conceals exactly why the learning is complete, and whether this is simply an artifact of the continuum players assumption.

2. THE EXAMPLE

2.1 Private Information and Prior Beliefs

An infinite sequence of individuals $n = 1, 2, \ldots$ sequentially takes actions in that exogenous order. Each individual can choose from one of two actions.

We first introduce a background probability space $(\Omega, \mathcal{E}, \nu)$. This space underlies all random processes in the model, and is assumed to be common knowledge.

There are two *states of the world* (or more simply, *states*), labelled H ('high') and L ('low'). Formally, this means that the background state space Ω is partitioned into two events Ω^H and Ω^L, called H and L. Let the common prior belief be that $\nu(H) = \nu(L) = 1/2$. That individuals have common priors is a standard modelling assumption, see e.g. Harsanyi (1967–68). Also, flat priors are truly WLOG, for SS show that more general priors are formally equivalent to a renormalization of the payoffs. Individual n receives a private random signal, $\sigma_n \in \Sigma$, about the state of the world.

Conditional on the state, the signals are assumed to be i.i.d. across individuals. It is common knowledge that in state H (resp. state L), the signal is distributed according to the probability measure μ^H (resp. μ^L). In the specific example that we consider, μ^H and μ^L are described by densities on $(0, 1)$, μ^H has the density $g(\sigma) = (1 - \sqrt{\sigma})/\sqrt{\sigma}$ while μ^L has the density 1. Following his observation of σ, the individual applies Bayes' rule to deduce what we shall refer to as his *private belief* $p(\sigma) = g(\sigma)/(g(\sigma) + 1) = 1 - \sqrt{\sigma} \in (0, 1)$ that the state is H. Conditional on the state, private beliefs are i.i.d. across individuals because signals are. In state H (resp. state L), p is distributed with a c.d.f. F^H (resp. c.d.f. F^L) on $(0, 1)$.

This may look very capricious, but there is method to our madness. Imagine that individuals are told one of two possible statistically true statements "with chance p, the state is high/low", where the signal quality p is uniform on $(0, 1)$. After being told that the state is high (resp. low), the individual's posterior equals p (resp. $1 - p$), since his prior was flat. Thus, in the high state, the density of individuals with posterior belief p is $f^H(p) = p/[p + (1 - p)] = p$, yielding $F^H(p) = p^2$. Similarly, in the low state, $F^L(p) = 2p - p^2$, for all $p \in (0, 1)$. This story of binary signals (i.e. being told 'high' or 'low') with variable certitude is less general than an arbitrary state-dependent signal with support Σ, since it implies the symmetry property that we exploit below.

We see that in the terminology of SS, the beliefs are unbounded, i.e. the support of F^H and F^L is all of $[0, 1]$. The intent with this example is to show that the complete learning result from SS carries over to this framework, even though the 'publicly' observable information is not refined over time.

2.2 The Dynamic Behavior of Posterior Beliefs

We assume WLOG that action 1 (eg., Invest) has payoff 1 in state H, and payoff -1 in state L, while action 2 is a neutral decision (eg., Don't Invest) having payoff 0 in each of the two states. Before deciding upon an action, the individual combines his *sample information* with his private signal, and arrives at his private posterior belief r that the state is H. The expected value of action 1 is thus $r - (1 - r) = 2r - 1$, and so the individual

optimally picks action 1 exactly when $r \geq 1/2$.

There is really an inductive aspect to the derivation of the posterior beliefs. We could first define the first individual's decision rule as a function of his private belief, and then describe how the second individual bases his decision on his private belief *and* on the observation of the first individual's action, and so on. Figures 1 and 2 illustrate the calculated values of the first few probabilities and thresholds along this line. But instead we shall go right in and simply say that individual n knows the decision rules of all the previous individuals, and bases his own decision on that knowledge.

Figure 1: **Private Belief Thresholds.** This portrays all possible values of the threshold $\bar{p}_n(k)$ for the first four individuals. Here, "Invest" is action 1, and "Don't" is action 2.
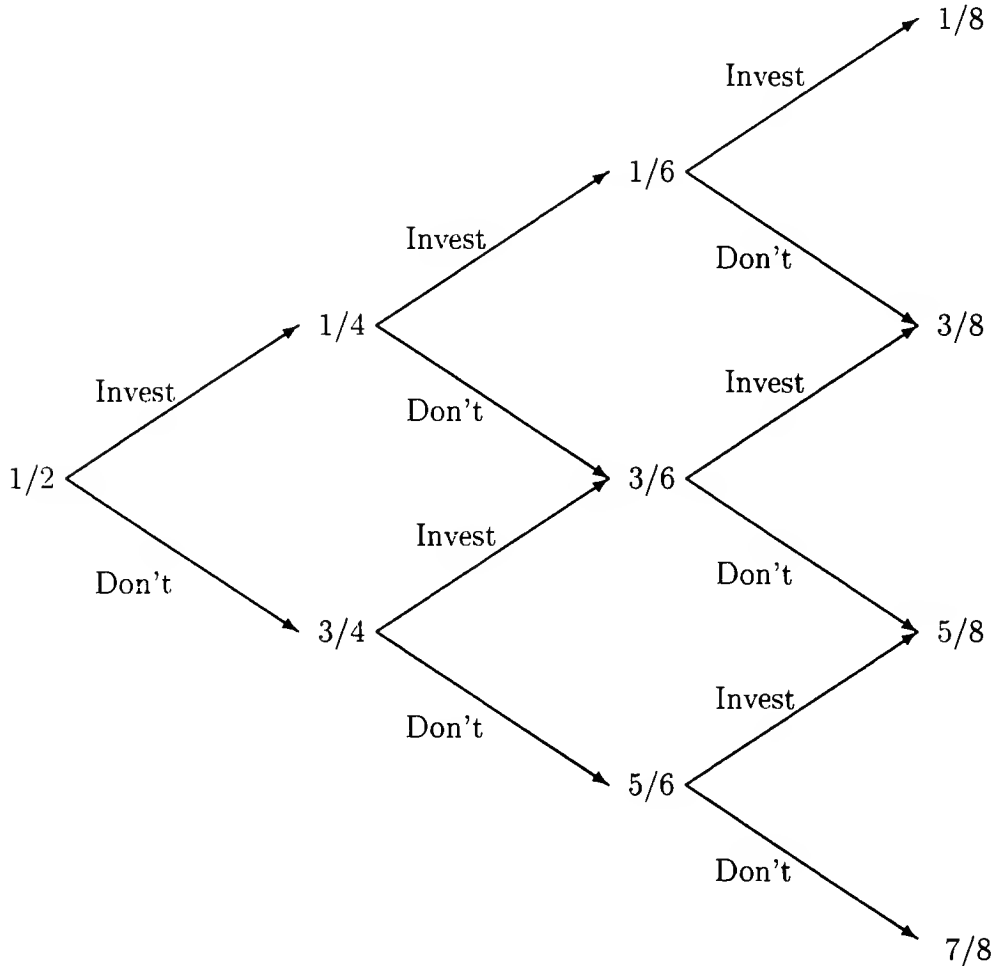
The sample information privy to individual n can be conveniently summarized by the number of predecessors who took action 1. More formally, we shall refer to $h_n(k)$ as the *history* (observed solely by individual n) that k of the first $n-1$ agents took action 1 (i.e. invested). Individual n can then calculate the various possible permutations of all possible action sequence. As individual n knows how his predecessors' decision rule as a function of their private signals and the history they could observe, he is able to calculate the

probability of each of their actions in either of the two states, under each of the possible action sequence permutations. This enables him to calculate the ex ante chance of any observed history $h_n(k)$ in either of the two states; we shall denote these probabilities by $\pi_n^H(k)$ and $\pi_n^L(k)$ (see figure 2). Finally, individual n is able to form a likelihood ratio that the state is low, namely $\ell_n(k) = \pi_n^L(k)/\pi_n^H(k)$. When k is not specified, we shall refer to this likelihood ratio as ℓ_n.

Figure 2: **Probabilistic Evolution.** This illustrates (at the tips of the arrows) the probabilities $\pi_n^H(k)$ for all initial histories up to period 3. Over the arrows are the transition probabilities in state H, i.e. $1 - F^H(\bar{p})$ of "Invest" and $F^H(\bar{p})$ of "Don't", by equation (2). Below the arrows are the corresponding transition probabilities in state L, respectively $F^L(\bar{p})$ and $1 - F^L(\bar{p})$.
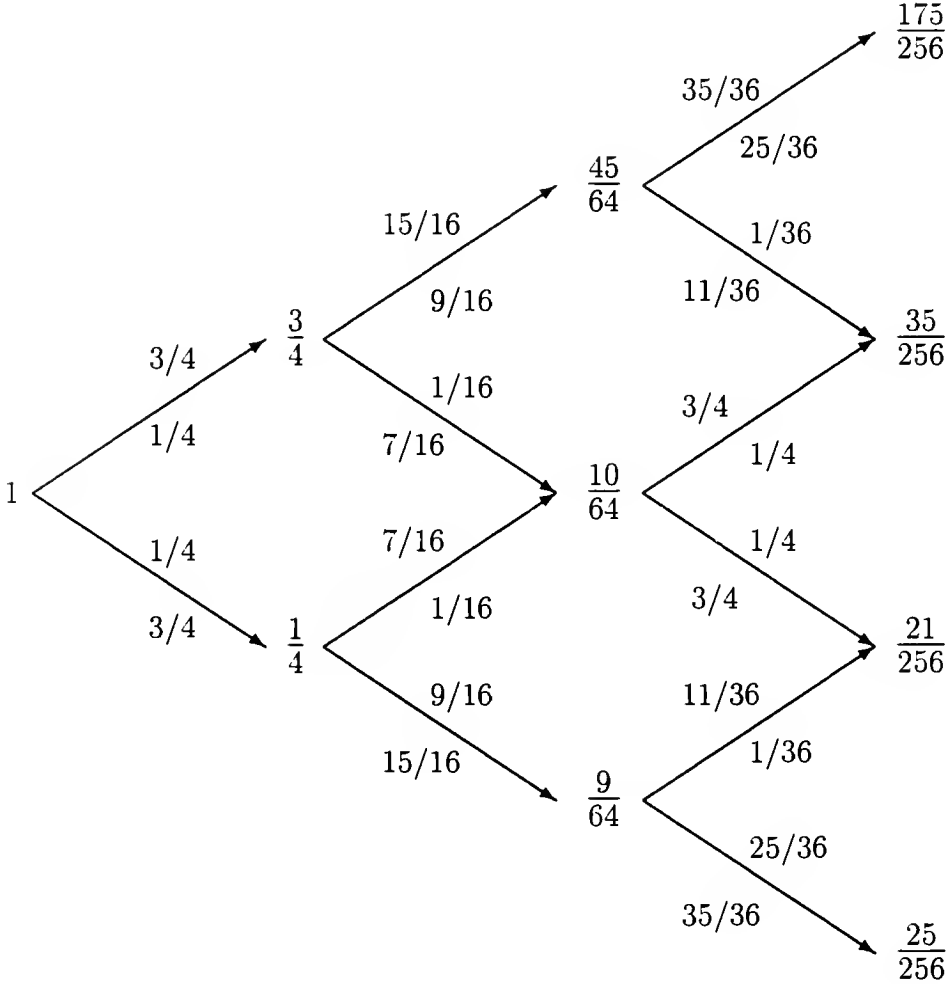
The first individual only observes the null history, and so we shall normalize $\ell_1(0) = 1$; later on, this will start a simple inductive derivation of the likelihood ratios. As private signals, and thereby individual actions, are random, $\langle \ell_n \rangle_{n=1}^{\infty}$ is a stochastic process. Using the likelihood ratio and the private belief, individual n may form his posterior belief using

Bayes' rule. In state H, this is given by

$$r = \frac{p\,\pi_n^H(k)}{p\,\pi_n^H(k) + (1-p)\,\pi_n^L(k)} = \frac{p}{p + (1-p)\,\ell_n(k)} \tag{1}$$

Observe that this expression is strictly increasing in p. We may interpret (1) as simply saying that the posterior odds $(1-r)/r$ of state L equal the likelihood ratio times the private odds. We see that $r \geq 1/2$ exactly when

$$p \geq \frac{\ell_n(k)}{1 + \ell_n(k)} = \frac{\pi_n^L(k)}{\pi_n^L(k) + \pi_n^H(k)} \equiv \bar{p}_n(k) \tag{2}$$

We shall call $\bar{p}_n$ the *public beliefs* in period n, since this would be individual n's posterior belief (that the state is *low*) had he a purely neutral private belief.[3] Individual n's decision rule is obvious: He takes action 1 exactly when his private belief p_n exceeds $\bar{p}_n$, or equivalently when his posterior belief r_n exceeds $1/2$.

We are interested in the asymptotic behavior of beliefs as expressed by the sequence of likelihood ratios $\langle \ell_n \rangle$, or public beliefs $\langle \bar{p}_n \rangle$. While in SS the likelihood ratios constitute a martingale, this is not true here. Later on, we explicitly reject this and other reasonable candidates for martingales. We shall, however, exploit the fact that $\langle \bar{p}_n \rangle$ is a supermartingale conditional on the state.

Particular to the information structure of this example, but *not* a general property of the observational learning paradigm, is a key symmetry property: $F^H(1-p) = 1 - F^L(p)$. Even though the closed form solutions will repeatedly exploit this fact, we are confident that it is not the driving force behind the complete learning. An implication of the symmetric posterior distributions are symmetric private belief thresholds, and a symmetric probabilistic evolution.

Lemma 1 (Symmetry) *For all n and $k = 0, \ldots, n-1$, we have $\bar{p}_n(k) = 1 - \bar{p}_n(n-1-k)$ and $\pi_n^H(k) = \pi_n^L(n-1-k)$.*

Proof: We proceed by induction on n. First $\pi_1^H(0) = \pi_1^L(0) = 1$ and $\bar{p}_1(0) = 1/2$, and so both formulae initially obtain. Let $N > 1$, and assume that both equations hold for all $n < N$. Let us now establish the equations for period N.

Fix any sequence of choices by the first $N-1$ individuals. Individual N knows exactly how many of his predecessors invested, and considers all the possible orders of moves that could lead to this outcome. For a given such order, let $k(n)$ denote how many of the first $n \leq N-1$ individuals have invested, and $\bar{p}_n(k(n))$ the private belief threshold for individual $n \leq N-1$. Denote by A the set of individuals who took action 1 in this particular ordering and by B the set $\{1, \ldots, N-1\} \setminus A$. Then the inductive hypothesis and $F^H(1-p) = 1 - F^L(p)$ yield

$$\prod_{n \in A} [1 - F^H(\bar{p}_n(k(n)))] \prod_{m \in B} F^H(\bar{p}_m(k(m)))$$

---

[3]This label, taken from SS, is really rather inappropriate here, since only individual n ever sees the public belief $\bar{p}_n$.

$$= \prod_{n \in A} F^L(1 - \bar{p}_n(k(n))) \prod_{m \in B} [1 - F^L(1 - \bar{p}_m(k(m)))]$$

$$= \prod_{n \in A} F^L(\bar{p}_n(n - 1 - k(n))) \prod_{m \in B} [1 - F^L(\bar{p}_m(m - 1 - k(m)))]$$

When individual N sums the above probabilities corresponding to all the possible sequences that could lead to the outcome at hand, he will clearly discover that $\pi_N^H(k) = \pi_N^L(N - 1 - k)$. We can thus conclude

$$\ell_N(k) = \frac{\pi_N^L(k)}{\pi_N^H(k)} = \frac{\pi_N^H(N - 1 - k)}{\pi_N^L(N - 1 - k)} = 1/\ell_N(N - 1 - k)$$

Application of (2) then yields $\bar{p}_n(k) = 1 - \bar{p}_n(n - 1 - k)$. $\diamond$

Lemmas 1 and 2 easily imply the relations

$$\bar{p}_n(k) = \frac{\pi_n^L(k)}{\pi_n^L(k) + \pi_n^H(k)} = \frac{\pi_n^H(n - 1 - k)}{\pi_n^H(n - 1 - k) + \pi_n^H(k)} \tag{3}$$

$$\ell_n(k) = \frac{\pi_n^H(n - 1 - k)}{\pi_n^H(k)} = \frac{\bar{p}_n(k)}{1 - \bar{p}_n(k)} \tag{4}$$

We now turn to the key difference equation that describes the evolution of the system over time. When individual n is told that k predecessors took action 1, he may conclude that one of two scenarios confronted his immediate predecessor: Either, individual $n - 1$ was told that $k - 1$ predecessors had taken action 1, and he then took action 2, or individual $n - 1$ was told that k predecessors had taken action 1, and he then took action 2. This leads to, for $1 \leq k \leq n - 2$,

$$\pi_n^H(k) = \pi_{n-1}^H(k - 1) \left[1 - F^H(\bar{p}_{n-1}(k - 1)) \right] + \pi_{n-1}^H(k) F^H(\bar{p}_{n-1}(k)) \tag{5}$$

One could assert the corresponding relationship in state L, but by Lemma 1 that is already embodied in (5). The relation also vacuously holds when $k = 0$ or $k = n - 1$, for then one of the terms in the sum vanishes. It is equation (5) that drives the dynamics of the model. In principle, one can always use it to calculate all probabilities (and likelihoods and thresholds) for the model.

2.3 Closed Form Solutions

We now summarize a slough of calculations, whose derivations are relegated to the Appendix (which non-skeptical readers should feel free to skip). The result describes the probabilistic evolution of the model in state H, and the public beliefs that arise after every history.

Lemma 2 *For all $n \geq 1$, $k = 0, \ldots, n - 1$ we have*

$$\bar{p}_n(k) = \frac{2(n - k) - 1}{2n} \tag{6}$$

$$\pi_n^H(k) = \frac{2n - 1}{2(n - k) - 1} 2^{4 - 4n} \binom{2k}{k} \binom{2(n - 1 - k)}{n - 1 - k} \binom{2(n - 1)}{n - 1} \tag{7}$$

6

One can verify that the numbers in figures 1 and 2 agree these formulae.

One nice property of the example is that we can explicitly the probability that everyone chooses the correct action from the very outset. For the chance that the first n individuals take the correct action 1 in state H is

$$
\begin{aligned}
\pi_{n+1}^H(n) &= (2n+1)2^{-4n}\binom{2n}{n}^2 \\
&= (2n+1)\left[\frac{(2n)!}{2^{2n}(n!)^2}\right]^2 \\
&= (2n+1)\left[\frac{(2n-1)(2n-3)\cdots 3\cdot 1}{2n(2n-2)\cdots 4\cdot 2}\right]^2 \\
&= \frac{[(2n+1)(2n-1)]\,[(2(n-1)+1)(2(n-1)-1)]\cdots[5\cdot 3]\,[3\cdot 1]}{(2n\cdot 2n)(2(n-1)\cdot 2(n-1))\cdots(4\cdot 4)(2\cdot 2)} \\
&= \prod_{j=1}^{n}\frac{4j^2-1}{4j^2}
\end{aligned}
$$

By Wallis' formula, the limit of this expression is

$$
\prod_{j=1}^{\infty}\frac{4j^2-1}{4j^2}=\frac{2}{\pi}
$$

where $\pi=3.14159\ldots$ is the well-known constant. That is, *the chance that all individuals invest in state H is $2/\pi$!* This result stands in contrast to the fact that there is zero chance that all individuals will take the wrong action in state H.

3. CONVERGENCE

Unfortunately, the discrete distribution found in Lemma 2 is unknown (to us!) despite its simple form and similarity to the hypergeometric distribution. But we can calculate its mean and prove that it converges to 0. This is the desired limit, for a zero threshold implies that all individuals invest, and that is indeed the correct action in state H.

Lemma 3 (Convergence in Mean) *In state H, the stochastic process $\langle \bar{p}_n \rangle$ converges to 0 in mean.*

Proof: For any $n \geq 2$, the mean of $\bar{p}_n$ satisfies[4]

$$
\begin{aligned}
E\bar{p}_n &= \sum_{k=0}^{n-1}\pi_n^H(k)\bar{p}_n(k) = \frac{2n-1}{2n}2^{4-4n}\binom{2(n-1)}{n-1}\sum_{k=0}^{n-1}\binom{2k}{k}\binom{2(n-1-k)}{n-1-k} \\
&= \frac{2n-1}{2n}2^{4-4n}\binom{2(n-1)}{n-1}4^{n-1} = \frac{2n-1}{2n}2^{2-2n}\frac{(2n-2)!}{(n-1)!\,(n-1)!} \\
&= 2^{1-2n}\frac{(2n-1)!}{n!\,(n-1)!} = \frac{1}{2}\frac{(2n-1)(2n-3)\cdots 3\cdot 1}{2n(2n-2)\cdots 4\cdot 2}
\end{aligned}
$$

---

[4]Below, we exploit the combinatorial identity $4^n = \sum_{k=0}^{n-1}\binom{2k}{k}\binom{2(n-1-k)}{n-1-k}$, for which thank several individuals on the internet newsgroup sci.math.

It is not hard to prove that this latter expression converges to 0. One way of doing so is as follows: As $\langle E\bar{p}_n \rangle$ is a decreasing sequence, and bounded below by 0. it converges to some nonnegative limit. If its limit were strictly positive, then we could take its logarithm:

$$\log\left[\lim_{n\to\infty} \frac{(2n-1)(2n-3)\cdots 3\cdot 1}{2n(2n-2)\cdots 4\cdot 2}\right] + \log(\tfrac{1}{2}) = \sum_{i=0}^{\infty} \log\left[\frac{2i+1}{2i+2}\right] - \log 2$$

$$\leq \sum_{i=0}^{\infty}\left[\frac{2i+1}{2i+2} - 1\right] = \sum_{i=0}^{\infty}\left[\frac{-1}{2i+2}\right] = -\infty$$

Hence, $\bar{p}_n \to 0$ in mean. $\diamond$

By equations (1) and (2), we have

$$r_n = \frac{p(1-\bar{p}_n)}{p(1-\bar{p}_n) + \bar{p}_n(1-p)}$$

Thus, by Lemma 3, the posterior belief sequence $\langle r_n \rangle$ converges to 0 in probability by Chebyshev's inequality. Still, we cannot yet conclude almost sure convergence of $\langle r_n \rangle$.[5] For this, we think that we must use the Martingale Convergence Theorem. Figure 3 provides a flavor of the nature of the convergence here.

Figure 3: **Empirical Density Function.**

Probability of various   histories for individual 100

---
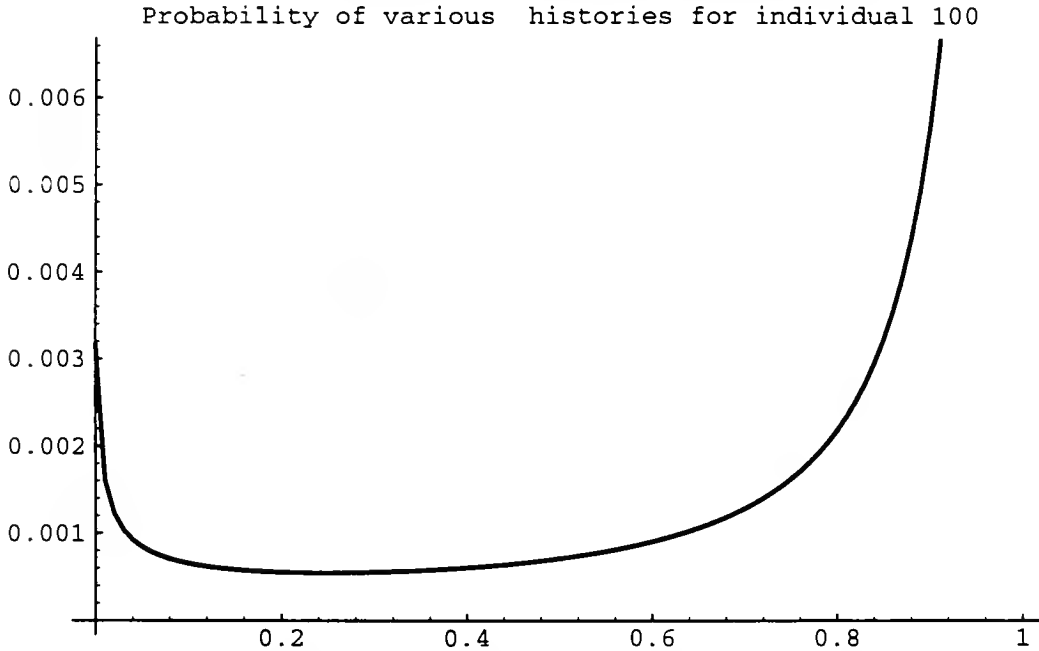
[5]To recall why this is so, consider the following well-known example (which appears for instance as Example 7.4 in Bartle (1966)). Let $\langle X_n \rangle_{n=1}^{\infty}$ be a sequence of stochastic variables on the probability space $[0,1]$ with Lebesgue measure. To any given $n \geq 1$ choose $a(n)$ such that $2^{a(n)} \leq n < 2^{a(n)+1}$. Let X_n be the indicator function of the subinterval $[\{n - 2^{a(n)}\}2^{-a(n)}, \{n+1-2^{a(n)}\}2^{-a(n)}] \subset [0,1]$. Clearly, $0 \leq X_n \leq 1$, and $EX_n = 2^{-a(n)} \to 0$; thus, $X_n \to 0$ in mean and probability. But, there is no $\omega \in [0,1]$ for which $X_n(\omega) \to 0$, because the value 1 will pop up infinitely often.

Consider the set of all possible infinite histories, and treat a finite history as the union of all infinite histories that contain it. Endow the set with the σ-algebra $\mathcal{A}$ that is generated by all possible finite length histories. All infinite length histories must be in $\mathcal{A}$ as well, because an infinite history is the countable intersection of finite histories. The probability measure on $\mathcal{A}$ is that of $\bar{p}_n$, so that history $h_n(k)$ has probability $\pi_n^H(k)$.

Now consider the sequence of sub-σ-algebras $\langle \mathcal{A}_n \rangle_{n=1}^\infty$, where $\mathcal{A}_n$ is generated by all finite histories of length n. Note that $\langle \mathcal{A}_n \rangle_{n=1}^\infty$ is an increasing sequence of σ-algebras because any history of length $n-1$ is the union of the two histories of length n that can follow it. Also, $\bar{p}_n$ is clearly measurable w.r.t. $\mathcal{A}_n$. The justification for introducing $\langle \mathcal{A}_n \rangle$ is that we can show that the process $\langle \bar{p}_n, \mathcal{A}_n \rangle$ is a supermartingale, which in turn establishes the almost sure convergence.

Lemma 4 (Conditional Supermartingale) $\langle \bar{p}_n, \mathcal{A}_n \rangle$ *is a supermartingale in state H.*

Proof: We start by verifying an equation which mysteriously holds in this example.

$$\begin{aligned}
\bar{p}_n(k) &= (1/2)\left[F^H(\bar{p}_n(k)) + F^L(\bar{p}_n(k))\right]\bar{p}_{n+1}(k) \\
&\quad + (1/2)\left[\left(1 - F^H(\bar{p}_n(k))\right) + \left(1 - F^L(\bar{p}_n(k))\right)\right]\bar{p}_{n+1}(k+1)
\end{aligned} \tag{8}$$

for all n and $k = 0, \ldots, n-1$. We do not claim that (8) holds beyond the current specification, as we in fact can show by counterexample that it does not. But its validity here is easily verified it by appealing to the closed form solutions of Lemma 2. For instance, rewrite (8) as

$$\bar{p}_n(k) - \bar{p}_{n+1}(k+1) = (1/2)\left[F^H(\bar{p}_n(k)) + F^L(\bar{p}_n(k))\right](\bar{p}_{n+1}(k) - \bar{p}_{n+1}(k+1))$$

But by (6) and $F^H(p) + F^L(p) = p^2 + (2p - p^2) = 2p$, this is equivalent to

$$(2(n-k) - 1)\left(\frac{1}{2n} - \frac{1}{2n+2}\right) = \frac{2(n-k) - 1}{2n}\frac{2}{2n+2}$$

which certainly holds.

We now prove that $E[\bar{p}_{n+1} \mid H, \bar{p}_n = \bar{p}_n(k)] < \bar{p}_n(k)$, which is the supermartingale inequality for state H. Just rewrite (8) as

$$\begin{aligned}
\bar{p}_n(k) &= \frac{1}{2}\left[F^H(\bar{p}_n(k))\bar{p}_{n+1}(k) + \left(1 - F^H(\bar{p}_n(k))\right)\bar{p}_{n+1}(k+1)\right] \\
&\quad + \frac{1}{2}\left[F^L(\bar{p}_n(k))\bar{p}_{n+1}(k) + \left(1 - F^L(\bar{p}_n(k))\right)\bar{p}_{n+1}(k+1)\right]
\end{aligned}$$

From the closed form solutions we know that $\bar{p}_{n+1}(k) > \bar{p}_{n+1}(k+1)$. We also know that F^L is always larger than F^H. Therefore, the latter of the two summands is larger. Since the two terms average out to give $\bar{p}_n(k)$, the first term must be smaller than $\bar{p}_n(k)$, i.e.

$$F^H(\bar{p}_n(k))\bar{p}_{n+1}(k) + (1 - F^H(\bar{p}_n(k)))\bar{p}_{n+1}(k+1) < \bar{p}_n(k)$$

This is exactly the statement that the supermartingale inequality holds in state H. $\diamond$

Finally, we arrive at our stated goal.

Proposition 1 (Complete Learning) *Almost surely, $\ell_n \to 0$ in state H, and $1/\ell_n \to 0$ in state L.*

Proof: By Lemma 4, $\langle \bar{p}_n \rangle$ is a supermartingale in state H. The Martingale Convergence Theorem therefore assures us that it converges a.s. to a r.v. $\bar{p}_\infty$. As $\langle \bar{p}_n \rangle$ is bounded, nonnegative, and converges in mean to 0, we know that a.s. $\bar{p}_\infty = 0$. So (4) implies $\ell_n \to 0$. $\diamondsuit$

It is worth remarking that complete learning here does not imply that eventually all individuals must take the correct action; or, in the terminology of SS, we need not eventually enter a 'correct herd'. The basic reason is the failure of what we called the 'overturning principle'. Suppose that one million individuals have taken action 1, and then the next individual takes action 2. That will hardly affect the likelihood ratio. Indeed, when individual one million and two observes his summary statistic of history (that all but one took action 1), he will find most likely the hypothesis that the individual who took the contrary action was one of the early ones; therefore, the unexpected action which bespeaks a signal more powerful than the entire weight of history, will have a negligible effect, simply because the information is not refined. In SS, we explain why we cannot rule out an infinite string of such isolated contrary actions — even whilst the likelihood ratio converges to 0.

4. OTHER MARTINGALES?

Finally, and partially in response to the stated disbelief of some, we wish to show that all other obvious (and nontrivial!) candidates for (sub-)martingales in the model fail the test. In particular, we shall consider the public beliefs and the likelihood ratios. Note that the private beliefs $\langle p_n \rangle$ clearly cannot be a conditional (sub-)martingale because they are i.i.d., and cannot be an unconditional martingale because obviously $\|\partial E[p_{n+1} \,|\, p_n]/(\partial p_n)\| < 1$. Similarly, the realized posterior beliefs $\langle \bar{r}_n \rangle$ could not possibly be a (sub-)martingale. Finally, it is easy to show that the likelihoods $\langle \pi_n^H \rangle$ is not a (sub-)martingale.

The analysis below uses the calculated thresholds and probabilities from figures 1 and 2.

1. PUBLIC BELIEFS.

Let us start with the sequence of public beliefs, namely $\langle \bar{p}_n \rangle$. In a single person experimentation or rational expectations framework (see, for instance, Aghion, Bolton, Harris and Jullien (1991) and Bray and Kreps (1987)), as well as the standard observational learning paradigm analyzed in SS and the herding literature, this stochastic process is a martingale unconditional on the state of the world. But figures 1 and 2 and Lemma 2 imply that

$$
\begin{aligned}
\bar{p}_1(1) &= 1/4 < 84/384 = (54/64)(1/6) + (10/64)(1/2) \\
&= [(3/4)(15/16) + (1/4)(9/16)](1/6) + [(3/4)(1/16) + (1/4)(7/16)](1/2) \\
&= \{[1 - \bar{p}_1(1)][1 - F^H(\bar{p}_1)] + \bar{p}_1(1)[1 - F^L(\bar{p}_1)]\}\bar{p}_2(2) \\
&\quad + \{[1 - \bar{p}_1(1)]F^H(\bar{p}_1) + \bar{p}_1(1)F^L(\bar{p}_1)\}\bar{p}_2(1) \\
&= Pr(\#2 \text{ invests} \mid \#1 \text{ invested}) \, \bar{p}_2(2) + Pr(\#2 \text{ doesn't} \mid \#1 \text{ invested}) \, \bar{p}_2(1) \\
&= E[\bar{p}_2 \mid k(1) = 1]
\end{aligned}
$$

Hence, the public beliefs are not an unconditional martingale. If one considers the history where #1 has moved and not invested, one can show that public beliefs are also not an unconditional supermartingale.

2. LIKELIHOOD RATIOS.

Next consider the likelihood ratio, which in SS is a martingale conditional on state H. We shall start at the very same history as above, where ℓ has the value $1/3$. In state H there is a probability of $15/16$ that its next value is $1/5$, and a probability of $1/16$ that its next value is 1. In expectation we get $1/4 < 1/3$. Similarly, by starting with #1 investing, we can show that it is not a supermartingale. Finally, unconditional on the state, it is easy to rule out a (sub-)martingale.

A. OMITTED PROOF OF LEMMA 2

Lemma 2 (Closed Form Solutions) *For all $n \geq 1$, $k = 0, \ldots, n - 1$ we have*

$$\bar{p}_n(k) = \frac{2(n - k) - 1}{2n}$$

$$\pi_n^H(k) = \frac{2n - 1}{2(n - k) - 1} 2^{4 - 4n} \binom{2k}{k} \binom{2(n - 1 - k)}{n - 1 - k} \binom{2(n - 1)}{n - 1}$$

Proof: CASE 1: $k = 0$ OR $k = n - 1$.

This is particularly simple because everyone can exactly deduce what all predecessors knew and did. Consider first $k = 0$. We shall proceed by induction on n. It is readily seen that the formulas are correct for $n = 1$: namely, $\bar{p}_1(0) = 1/2$ and $\pi_1^H(0) = 1$. Now assume that it holds for arbitrary $n \geq 1$, and we shall prove it holds for $n + 1$ (still with $k = 0$).

Individual $n + 1$ is thus told that none of the first n individuals invested; therefore, he can then conclude that individual n had been told that none of the first $n - 1$ individuals invested. So $\bar{p}_n = (2n - 1)/(2n)$, and by (4), $\ell_n = 2n - 1$. Individual n took action 2. From the viewpoint of individual $n + 1$, there was in state H a probability of $F^H(\bar{p}_n)$ that this could happen, and in state L a probability of $F^L(\bar{p}_n)$. Thus the next likelihood ratio is

$$\begin{aligned}
\ell_{n+1} &= \ell_n F^L(\bar{p}_n)/F^H(\bar{p}_n) \\
&= \ell_n \left(\frac{2(2n - 1)}{2n} - \frac{(2n - 1)^2}{(2n)^2} \right) \Big/ \frac{(2n - 1)^2}{(2n)^2} \\
&= (2n - 1) \frac{(2n - 1)(2n + 1)}{4n^2} \frac{4n^2}{(2n - 1)^2} = 2n + 1
\end{aligned}$$

which by (2) yields $\bar{p}_{n+1} = [2(n + 1) - 1]/2(n + 1)$, as desired. Also, equation (5) and the inductive hypothesis yields

$$\begin{aligned}
\pi_{n+1}^H(0) &= \pi_n^H(0) F^H(\bar{p}_n(0)) \\
&= 2^{4 - 4n} \binom{2(n - 1)}{n - 1}^2 \frac{(2n - 1)^2}{(2n)^2} \\
&= 2^{4 - 4n} \frac{(2n - 2)!}{(n - 1)!(n - 1)!} \frac{(2n - 2)!}{(n - 1)!(n - 1)!} \frac{(2n - 1)^2}{(2n)^2} = 2^{-4n} \left[\frac{(2n)!}{n!n!} \right]^2
\end{aligned}$$

For $k = n - 1$ just apply Lemma 1. $\qquad \bullet$

11

CASE 2: $1 \leq k \leq n - 2$.

Again we shall use induction on n. The formula is true for $n = 2$, the first nonvacuous case, as $\bar{p}_3(1) = 1/2$ and $\pi_1^H(0) = 5/32$ follows from figures 1 and 2. Next we shall assume that the formulae hold for n, and establish them for $n + 1$.

We start by the induction for π^H. First rewrite (5) as

$$
\begin{aligned}
\pi_{n+1}^H(k) &= \pi_n^H(k - 1)\left[1 - F^H(\bar{p}_n(k - 1))\right] + \pi_n^H(k)F^H(\bar{p}_n(k)) \\[2mm]
&= \pi_n^H(k - 1)\left[1 - \left(\frac{2(n - k) + 1}{2n}\right)^2\right] + \pi_n^H(k)\left(\frac{2(n - k) - 1}{2n}\right)^2 \\[2mm]
&= \underbrace{\frac{-4n - 1 - 4k^2 + 8nk + 4k}{4n^2}\pi_n^H(k - 1)}_{term1} + \underbrace{\frac{(2(n - k) - 1)^2}{4n^2}\pi_n^H(k)}_{term2}
\end{aligned}
$$

Now, using the hypothesized formula for $\pi_n^H(k - 1)$, we can rewrite term 1 as

$$
\begin{aligned}
&\frac{(4n - 2k + 1)(2k - 1)}{4n^2}\pi_n^H(k - 1) \\[2mm]
&= \frac{(4n - 2k + 1)(2k - 1)(2n - 1)}{4n^2(2(n - k) + 1)2^{4n-4}}\binom{2k - 2}{k - 1}\binom{2(n - k)}{n - k}\binom{2n - 2}{n - 1} \\[2mm]
&= \frac{(4n - 2k + 1)k(2n - 1)}{4n^2(2(n - k) + 1)2^{4n-3}}\binom{2k}{k}\binom{2(n - k)}{n - k}\binom{2n - 2}{n - 1} \\[2mm]
&= \frac{(4n - 2k + 1)k}{n(2(n - k) + 1)2^{4n}}\binom{2k}{k}\binom{2(n - k)}{n - k}\binom{2n}{n} \\[2mm]
&= \frac{(4n - 2k + 1)k}{n(2n + 1)}\frac{(2(n + 1) - 1)}{(2(n + 1 - k) - 1)2^{4(n+1)-4}}\binom{2k}{k}\binom{2(n + 1 - 1 - k)}{n + 1 - 1 - k}\binom{2n}{n}
\end{aligned}
$$

and term 2 like this:

$$
\begin{aligned}
&\frac{(2(n - k) - 1)^2(2n - 1)}{4n^2(2(n - k) - 1)2^{4n-4}}\binom{2k}{k}\binom{2(n - 1 - k)}{n - 1 - k}\binom{2n - 2}{n - 1} \\[2mm]
&= \frac{n - k}{n2^{4n}}\binom{2k}{k}\binom{2(n - k)}{n - k}\binom{2n}{n} \\[2mm]
&= \frac{(n - k)(2(n - k) + 1)}{n(2n + 1)}\frac{(2(n + 1) - 1)}{(2(n + 1 - k) - 1)2^{4(n+1)-4}}\binom{2k}{k}\binom{2(n + 1 - 1 - k)}{n + 1 - 1 - k}\binom{2n}{n}
\end{aligned}
$$

Finally, adding terms 1 and 2 provides the desired expression for $\pi_{n+1}^H(k)$ because

$$
\frac{(4n - 2k + 1)k}{n(2n + 1)} + \frac{(n - k)(2(n - k) + 1)}{n(2n + 1)} = 1
$$

We now need only manipulate the formula for π_{n+1}^H and (3) to establish

$$\bar{p}_{n+1}(k) = \frac{\pi_{n+1}^H(n+1-k)}{\pi_{n+1}^H(n+1-k) + \pi_{n+1}^H(k)} = \frac{\frac{2n+1}{2k+1}}{\frac{2n+1}{2k+1} + \frac{2n+1}{2(n-k)+1}} = \frac{2(n+1-k)-1}{2(n+1)}$$

$\diamondsuit$

References

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien, "Optimal Learning by Experimentation," *Review of Economic Studies*, 1991, *58* (196), 621–654.

Banerjee, Abhijit, "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 1992, *107*, 797–817.

Bartle, Robert G., *The elements of integration*, New York: Wiley, 1966.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, 1992, *100*, 992–1026.

Bray, Margaret and David Kreps, "Rational Learning and Rational Expectations," in G. Feiwel, ed., *Kenneth Arrow and the Ascent of Economic Theory*, London: Macmillan, 1987, pp. 597–625.

Easley, David and Nicholas Kiefer, "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 1988, *56*, 1045–1064.

Harsanyi, J., "Games with Incomplete Information Played by Bayesian Players," *Management Science*, 1967–68, *14*, 159–182, 320–334, 486–502.

Smith, Lones, "Error Persistence, and Experiential versus Observational Learning," 1991. Foerder Series working paper.

— and Peter Sørensen, "Pathological Models of Observational Learning," 1994. MIT mimeo.

Date Due